

HP 12-Port 4X Fabric Copper Switch Installation Guide

HP-UX Server Networking

HP-UX11i v2



Manufacturing Part Number: AB291-96003

April 2004

Printed in the US

© Copyright 2004-2005 Hewlett-Packard Development Company L.P.

Legal Notices

The information in this document is subject to change without notice.

Hewlett-Packard makes no warranty of any kind with regard to this manual, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be held liable for errors contained herein or direct, indirect, special, incidental or consequential damages in connection with the furnishing, performance, or use of this material.

Warranty

A copy of the specific warranty terms applicable to your Hewlett-Packard product and replacement parts can be obtained from your local Sales and Service Office.

U.S. Government License

Proprietary computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Copyright Notice

Copyright © 1997-2004 Hewlett-Packard Development Company L.P. All rights reserved. Reproduction, adaptation, or translation of this document without prior written permission is prohibited, except as allowed under the copyright laws.

Trademark Notices

UNIX® is a registered trademark in the United States and other countries, licensed exclusively through The Open Group.

1 Installation Overview & Planning

NOTE The latest copy of this document can be found at <http://docs.hp.com> within the Networking and Communications area.

Overview of the HP Fabric Clustering System

This document contains the following sections that give installation information for HP Fabric:

- “About HP Fabric Clustering System for InfiniBand”
- “HP Fabric Clustering System Hardware Components”
- “HP Fabric Installation Prerequisites”
- “Install Fabric Adapters”
- “Install HP Fabric Software”
- “Install Fabric Switches”
- “Attach Cable Guides & Cables”
- “Internet Protocol over InfiniBand (IPoIB)”

About HP Fabric Clustering System for InfiniBand™

InfiniBand is an industry-standard high-speed, packet-based interconnect for node-to-node communications. HP Fabric Clustering System for InfiniBand provides higher speed, lower network latency and uses less CPU than other industry standard protocols (for example, Fibre Channel and Gigabit Ethernet). Instead of using a traditional bus-based technology, HP Fabric is built around switched fabric architecture, providing the bandwidth necessary for high speed data transfer. For more information on the InfiniBand standard see “<http://www.infinibandta.org>”.

HP’s release of this clustering solution delivers the performance and scalability required by the following:

- Technical Computing Clusters supporting HP Message Passing Interface (MPI) based applications

HP Fabric Clustering System Hardware Components

This section describes the various HP Fabric products. For more information on HP Integrity server systems that support HP Fabric products, see the *HP Fabric Clustering System Release Notes*, available at <http://www.docs.hp.com/en/netcom.html#Fabric%20Clustering%20System%20for%20InfiniBand>.

NOTE This document uses the term **HP Fabric Clustering System** to refer to the hardware and software that form the HP Fabric Clustering System interconnect product.

The term **HP Fabric Clustering System** refers to the following copper-based hardware components:

- AB286A HP PCI-X 2-port 4X Fabric (HPC) Adapter
- AB291A HP 12-port 4X Fabric Copper Switch

- AB346A HP 5m 4x Fabric Copper Cable
- AB353A HP 7m 4x Fabric Copper Cable
- AB347A HP 10m 4x Fabric Copper Cable

HP Fabric Clustering System Host Channel Adapter (HCA)

AB286A HP PCI-X 2-port 4X Fabric (HPC) Adapter

Figure 1-1 Host Channel Adapter



Fabric Switch

The AB291A HP 12-port 4X Fabric Copper Switch is as follows:

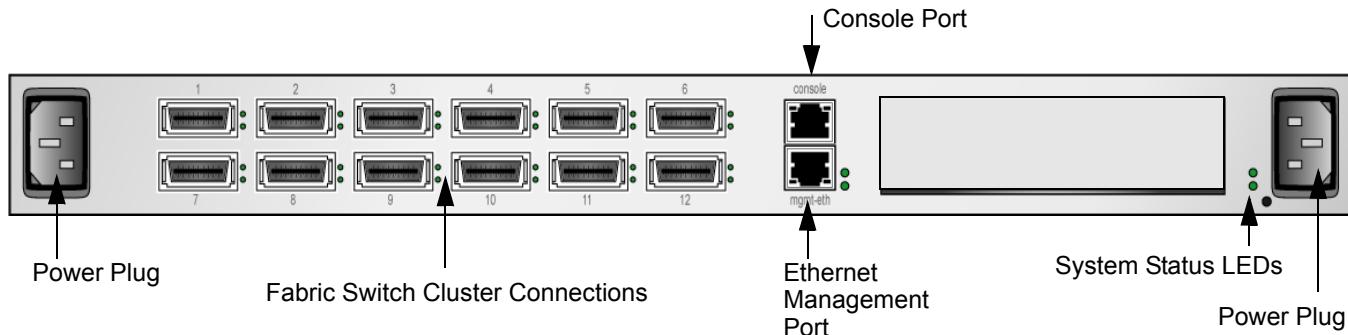
Figure 1-2 Fabric Switch Front



Figure 1-3 Rear View of Fabric Switch



Figure 1-4 Stylized View of Switch Connectors



Other Product Elements

Other elements of the HP Fabric product family:

- HP 4X Fabric Copper Cables
 - AB346A (5m length)
 - AB353A (7m length)
 - AB347A (10m length)
- The HP Fabric software consists of HP-UX drivers and user-space components as well as firmware on fabric adapters and switches.
- Fabric software supports Remote Direct Memory Access (RDMA). RDMA provides higher bandwidth, lower CPU overhead, and lower latency (the time a message takes to get from one point to another).

Installing HP Fabric Clustering System

The installation procedure for the HP Fabric clustering system installs the Switch, Host Channel Adapter (HCA), cables, cable guide, and gets the system up and running.

Managing the Switch

The fabric switch can be used without being managed. This Installation Section does not include information for managing HP Fabric. You can manage it using its Command Line Interface (CLI) - a text-based interface accessible through a direct serial connection, Telnet over IP, or SSH over IP.

HP Fabric Installation Prerequisites

Before installing the HP Fabric product, ensure that the following hardware and software prerequisites are met:

- Check the *HP Fabric Clustering System Release Notes* for known problems, required patches, or other information needed for installation, available at <http://www.docs.hp.com> in the Networking & Communications section.

- Confirm that the fabric adapter and software is supported on the HP Integrity server and IO slot chosen, see *HP Fabric Clustering System Support Matrix*, available at <http://www.docs.hp.com> in the Networking & Communications section.
- Determine where the switch is to be installed and which servers are to have an HCA installed.
- To use the Internet Protocol over InfiniBand (IPoIB) protocol, certain patches are required, along with installing the IPoIB software itself, see “Internet Protocol over InfiniBand” (IPoIB) on page 13 (optional).
- The installed version of IP Filter must be higher than “B9901AA A.03.05.06 HP IPFilter 3.5alpha5”.
- Check the cluster configuration plan for information on where to install adapters, switches, and cable guides, and where cables are to be routed.

For more information about the configuration plan, please refer to the *Planning the Cluster* section of the *HP Fabric Clustering System Support Guide*.

- Check the HP Integrity server’s documentation to determine if additional tools may be required for component installation.
- For specific instructions on adapter installation, see system-specific documentation on “installing networking adapters” for each type of HP Integrity server that fabric adapters will be installed into.
- On-Line Addition and Replacement (OL*) of HP Fabric adapters is supported on HP Fabric Clustering System version B.11.23.02 with a couple of limitations; see “On-Line Addition and Replacement Operations (OL*)” on page 14
- Confirm that there are cables of proper length to make connections in the fabric (adapter to adapter, adapter to switch, or switch to switch).
- HP Fabric software suite software media is available. It is included on CD/ or available on the web through <http://www.software.hp.com>.
- Confirm that HP-UX super-user privileges are available; they are necessary to complete the installation.
- Confirm that the /usr/bin, /usr/sbin, /sbin, and /opt/rdma/bin directories are in your PATH by logging in as root and using the echo \$PATH command.
- Confirm that the HP-UX operating system is the correct version, HP-UX 11i v2. Use the uname -a command to determine the HP-UX version of your system.

For more information about the required operating system versions, see *HP Fabric Clustering System Release Notes*.

Information Available on the Host Channel Adapter

If you are collecting information from the adapter before installation for inventory purposes, there is a GUID number on the card. This number can also be obtained through the HP-UX fabric utilities, but if recording information for inventory, include the GUID number as well, as it is unique for each adapter. The GUID will have a format such as: **00:05:ad:00:00:00:02:40**

Install Fabric Adapters

CAUTION	Fabric adapters contain electronic components that can easily be damaged by small amounts of electricity. To avoid damage, follow these guidelines:
	<ul style="list-style-type: none">• Store adapters in their anti-static plastic bags until installation.• Work in a static-free area, if possible.

- Handle adapters by the edges only. Do not touch electronic components or electrical traces.
- Use the disposable grounding wrist strap provided with each adapter. Follow the instructions included with the grounding strap.
- Use a suitable ground—any exposed metal surface on the computer chassis.

Installing the Fabric Host Channel Adapter

The adapter comes pre-configured. You do not have to set any jumpers or connectors.

The firmware for the HCA is compiled into the driver and the driver will download the current version onto the HCA if an older version is on the card.

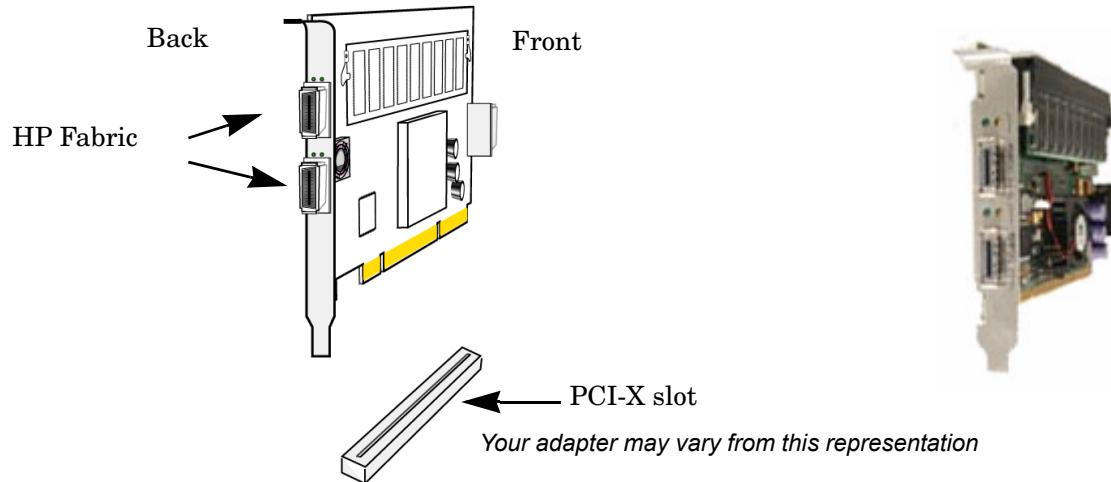
Step 1. Access the system card bay

- If the system is running, shut it down by executing: shutdown -h. Respond “y” to the continue to shutdown prompt.
- Wait for the system to shut down completely, and then power off the system by pressing the system off button. Ensure that the system is grounded.
- Open the system to gain access to the PCI backplane.
- Insert the card into the highest performance PCI-X slot available.

Step 2. Install the card

- Slide the adapter edge-connector into the PCI-X slot until it is fully seated.
- Secure the card and reassemble the system.

Figure 1-5 Inserting the Host Channel Adapter



Install HP Fabric Software

Install the HP Fabric Software

This section describes the steps necessary to install HP Fabric software. The software must be installed on each instance of the HP-UX operating system in the fabric.

Step 1. Log on to the system as root.

Step 2. Insert the software media into the appropriate drive. If the software is being loaded from a CD-ROM, go to step 3; otherwise, go to step 4. The software can also be located at <http://www.software.hp.com>.

Step 3. Mount the CD-ROM drive by using the following command at the command prompt:

```
$ mount device_name
```

where *device_name* is the name assigned to the CD-ROM drive.

Step 4. Run the *swinstall* program using the following command:

```
$ /usr/sbin/swinstall
```

This opens the **Software Selection** window.

Step 5. Change the Source Host Name, if necessary, and then enter the mount point of the drive in the Source Depot Path field. Select the **OK** button to return to the **Software Selection** window.

Step 6. The **Software Selection** window now contains a list of available software to install.

Step 7. Highlight the *HP Fabric for InfiniBand* software:

IB4X-00 -> B.11.23.01 PCI-X InfiniBand; Supptd HW=AB286A

Choose **Mark for Install** from the **Actions** menu; this chooses the highlighted software.

You can also install the IPoIB software at this time as it is part of this same bundle to enable it to be functional. It can also be installed later if you choose to not install it now.

Step 8. From the **Actions** menu, select the **Install...** menu, and then choose **Install**. This begins product installation and opens the **Install Analysis** window.

Step 9. Select the **OK** button in the **Install Analysis** window when the Status field displays a **Ready** message.

Step 10. Select the **YES** button in the **Confirmation** window to start software installation.

swinstall loads the fileset, runs the control script for the filesets, and builds the kernel. When the processing is finished, the **Status** field displays a **Ready** message. Select **Done** and then the **Note** window opens.

Step 11. Select the **OK** button in the **Note** window to reboot. The user interface disappears and the system reboots.

When the system comes back up, log on to the system as root and view the */var/adm/sw/swagent.log* and */var/adm/sw/swinstall.log* files to view error or warning messages that may have occurred during the installation.

Step 12. Verify that all installed fabric adapters have a software state of **CLAIMED**, by running the *ioscan -fk -C ib* command.

Install Fabric Switches

This section describes how to install the switch in a standard HP rack.

Requirements

In addition to the accessories provided with the switch, you should have:

- A T-10 Torx driver and a Phillips #2 screw driver.
- A second person for the installation.

Rack Mount Preparation

Prior to mounting the switch in a rack:

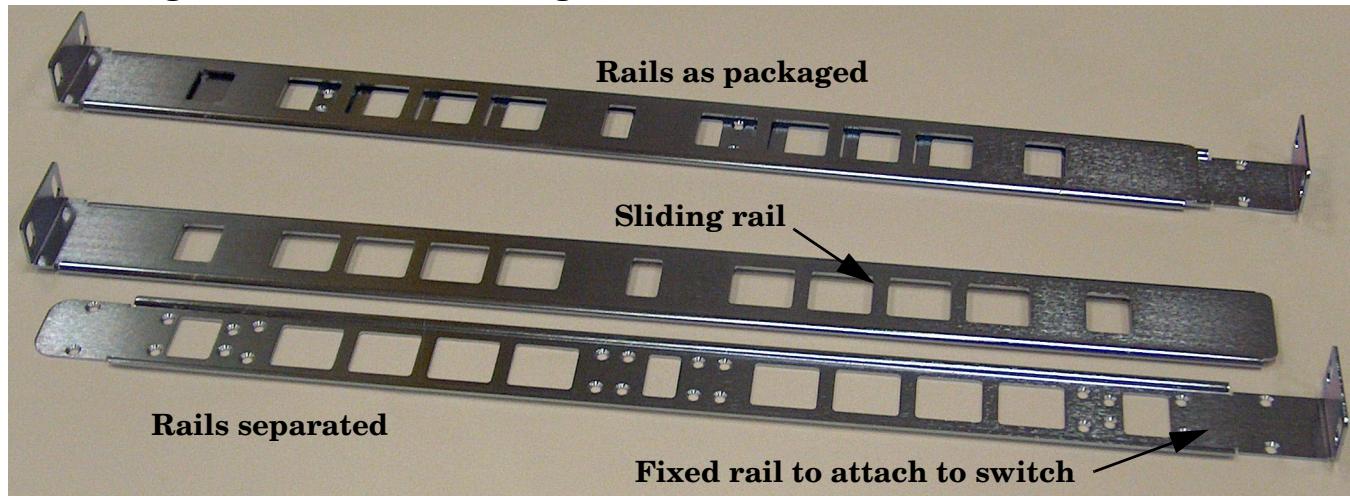
1. Check the slot in the rack for 1U clearance.
2. Open the box, remove the switch, rail brackets, and parts bag.
3. Place the switch on a secure, clean surface.
4. Open the plastic bag containing mounting parts and locate:
 - (12) Torx screws for attaching the rails to the switch
 - (8) screws and clips for attaching the rails to the rack.

NOTE It is highly recommend that two people install the switch, but it can be installed with one, however it may be difficult to support and align the switch correctly in the rack rails with only one person.

Installing the Switch

Step 1. Take one set of sliding brackets and separate them. Each side arrives with the sliding rail assembled to a fixed rail. They need to be separated prior to attaching to the switch.

Figure 1-6 Fixed and Sliding Rails



Step 2. To install the switch with the bezel to the front of the rack, install with the flange at the front of the switch. You can also mount the flange toward the back if you want to mount the switch backward in the rack such that the service-side is forward.

Attach the rail with 6 screws to each side of the switch in the locations shown in Figure 1-7

Screw locations are the same in the switch regardless of rail orientation.

Figure 1-7 Rail-to-Switch Attachment Locations

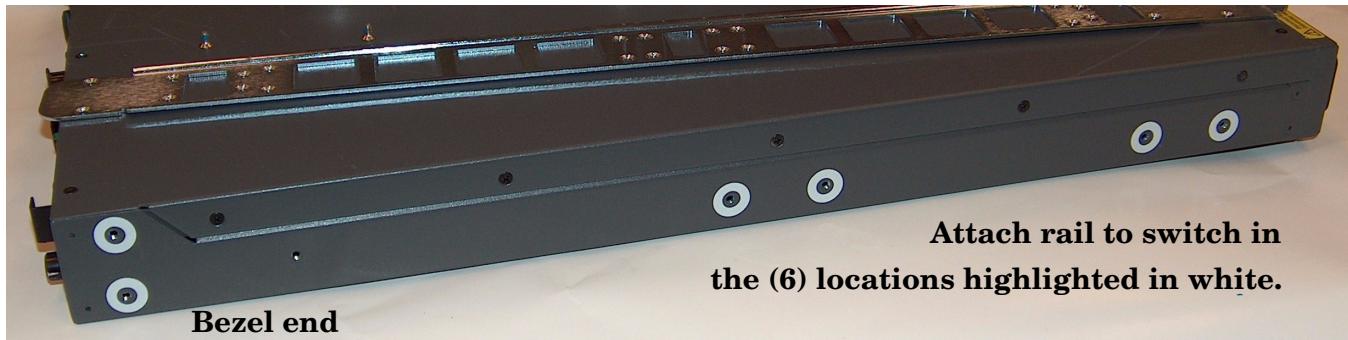
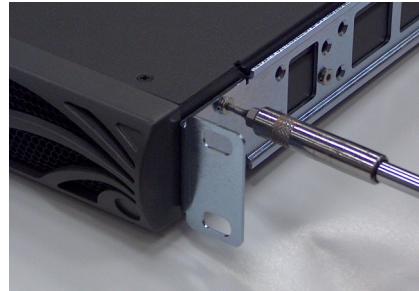


Figure 1-8 Attach Rail to Switch



Step 3. Repeat steps Steps 1 and 2 on the other side of the switch.

The two sliding rails should still be unattached.

Step 4. Loosely attach the switch with one screw through the front of each rail.

CAUTION Temporarily support the switch until the sliding rails are attached.

Figure 1-9 Loosely Attach Rail to Rack



CAUTION Do not let the switch hang unsupported from only the fixed rail flanges.

Step 5. Fit the sliding rails over the fixed rails with the flanges facing outward until the flanges meet the rack.

Figure 1-10 Fit Sliding Rail over Fixed Rail



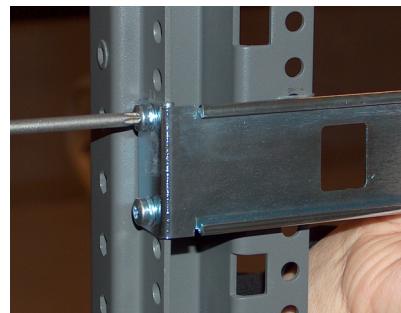
Figure 1-11 Fit Sliding Rail to Rack



Step 6. Attach the rails to the rack, both sides as shown in Figure 1-12.

CAUTION Keep switch supported until rail is attached to the rack.

Figure 1-12 Attach Sliding Rail to Rack



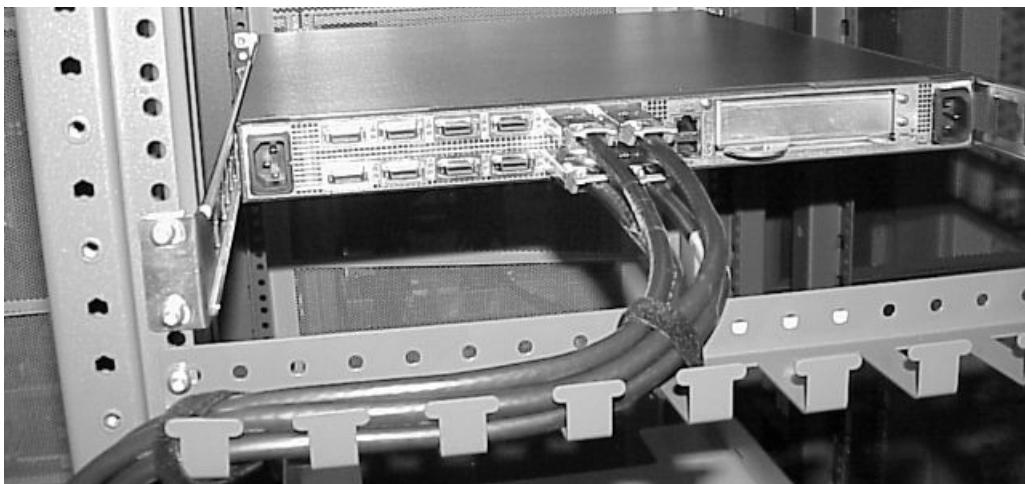
Step 7. Securely attach all rails to the rack.

Attach Cable Guides & Cables

The Fabric cables are high performance impedance-controlled cables with protective shielding, and due to the thickness and stiffness of the cables, HP recommends that Cable Guides be installed prior to attaching cables to minimize stress on the connectors for both switch and HCA cabling.

The Fabric cables should be securely attached to the rack enclosure. HP recommends the use of the included E7723-60003 Cable Guide Installation Kit mounted in a horizontal fashion as depicted below in Figure 1-13. A second Cable Guide is included and it can be installed in the horizontal or vertical position to better guide & support the cables for either the switch or HCA. Cable Guide Installation Instructions are provided as a separate document included with the Cable Guides.

Figure 1-13 **Cable Guide with Cables**



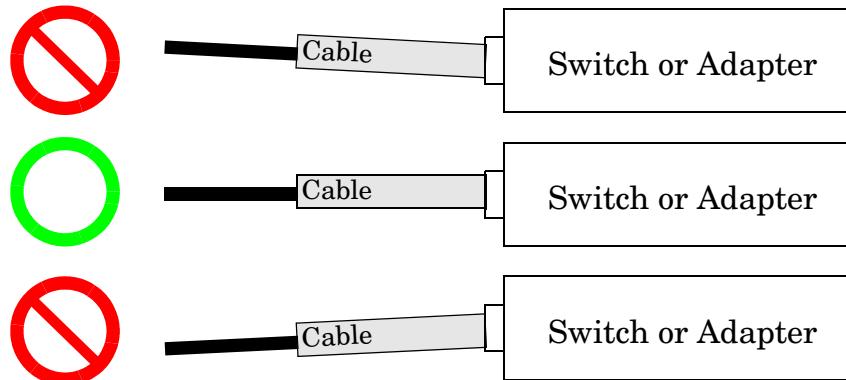
Attach Cables

The connectors include a latching mechanism. The latching and unlatching mechanism varies by manufacturer, and prior to installing cables, familiarize yourself with how to engage and disengage the latch.

Attach cables to switches and HCAs assuring that the latch is engaged on each connection.

Do not attempt to make sharp bends or folds in the cables. Care should be taken such that the mounting of the cable to switch or adapter does not place unnecessary stress on the connection as depicted below:

Figure 1-14 **Correct & Incorrect Cable Installation**



Attach to other Fabric devices:

1. Be sure you have attached a connector cable to the adapter. Push the connector in until you hear it click.
2. Attach the free end of the cable to a compatible fabric switch or adapter.

The following table lists the cables that you can use with HP Fabric.

Table 1-1 Cable Products

AB346A	5 meter 4X Fabric Copper Cable
AB353A	7 meter 4X Fabric Copper Cable
AB347A	10 meter 4X Fabric Copper Cable

Internet Protocol over InfiniBand™ (IPoIB)

In order to use the Internet Protocol over InfiniBand (IPoIB) the following patches need to be installed on each server used for IPoIB in the HP Fabric clustering system:

Patches can be downloaded from <http://www2.itrc.hp.com/service/patch/mainPage.do>

- Install DLPI patch, PHNE_30242
- Install TOUR product bundle (A.02.00 Transport Optional Upgrade Release for B.11.23)
The TOUR product can be found on <http://www.software.hp.com>.
- Install SAM patch, PHCO_29666
- If you did not include the IPoIB product during the initial HP Fabric installation, it needs to be installed at this time, see “Install the HP Fabric Software” on page 7

Enabling the IPoIB Interface

ifconfig Command

Description: This command enables or disables an IPoIB interface, configuring the link to the HP Fabric clustering system.

Syntax: ifconfig lan9000

Example: # lan9000

Output Example:

```
lan9000: flags=843<UP,BROADCAST,RUNNING,MULTICAST>
          inet 100.99.99.53 netmask ffffff800 broadcast 100.99.103.255
```

```
hptem052 1010: ifconfig lan9000 inet6
lan9000: flags=4800841<UP,RUNNING,MULTICAST,PRIVATE,ONLINK>
          inet6 fe80::202:c902:0:77d prefix 10
```

On-Line Addition and Replacement Operations (OL*)

On-Line Addition and Replacement (OL*) of HP Fabric adapters is supported on HP Fabric Clustering System version B.11.23.02 with a couple of limitations related to replacing an HCA, see below.

Replacing an HP Fabric Clustering System HCA

CAUTION Do Not use an Attention Button to REPLACE an HP Fabric HCA.

Do not use the Attention Button method of On-Line Addition and Replacement (OL*) to replace an existing HP Fabric Clustering System HCA, as the slot will not power off. Use the Command Line Interface, or Peripheral Device Tool graphical user interface (GUI) to replace an HP Fabric HCA. Refer to the *Interface Card OL* Support Guide*, available on <http://www.docs.hp.com> for details on using standard OL* features.

To replace an HCA using the command line, you must use option `olrad -f` as in this example:

```
$ /usr/bin/olrad [-f] -r slot_id
```

where `slot_id` is the Slot ID of the slot.

The `-f` option overrides the “data critical” errors returned by Critical Resource Analysis. It is important to note that `olrad` will not allow “critical” errors to be overridden and that `olrad` automatically overrides “warnings”.

Whether `-f` is specified or not, critical resource analysis routines are run before an OLA/R operation, to ensure that the current OLA/R operation does not interrupt the normal working of the system; in other words, to identify “critical” errors.

The “data critical” errors are typically not critical to the system, but they may be critical to the user. Hence, the user need to decide whether or not to use the `-f` option for overriding these types of errors.

HP Fabric HCAs can be added using the Attention Button feature.